arXiv:1411.3302v1 [cs.LG] 12 Nov 2014

Using Gaussian Measures for Efficient Constraint Based Clustering

Chandrima Sarkar,[1*] Atanu Roy,[1]
sarkar@cs.umn.edu, atanu@cs.umn.edu
[1] Department of Computer Science and Engineering, University of Minnesota, Twin Cities
*To whom correspondence should be addressed; E-mail: sarkar@cs.umn.edu

In this paper we present a novel iterative multiphase clustering technique for efficiently clustering high dimensional data points. For this purpose we implement clustering feature (CF) tree on a real data set and a Gaussian density distribution constraint on the resultant CF tree. The post processing by the application of Gaussian density distribution function on the micro-clusters leads to refinement of the previously formed clusters thus improving their quality. This algorithm also succeeds in overcoming the inherent drawbacks of conventional hierarchical methods of clustering like inability to undo the change made to the dendogram of the data points. Moreover, the constraint measure applied in the algorithm makes this clustering technique suitable for need driven data analysis. We provide veracity of our claim by evaluating our algorithm with other similar clustering algorithms.

Introduction

In this paper we examine clustering, which is an important category of data mining problem. Data mining is the process of finding hidden patterns and trends in databases and using that

information to do a variety of tasks such as finding association rules, clustering heterogeneous groups of information and build predictive models [9]. It can also be considered as an important tool for data segmentation, selection, exploration and building models using the vast data stores to discover previously unknown patterns in various domains such as healthcare [12] [21], [23], [20] [19], social media analysis [18],[24], [2], [16], [11], [15], [17], finances and various other domains [22]. From the past few decades, data mining has been used extensively in various areas of decision making and decision analysis by financial institutions, for credit scoring and fraud detection; marketers, for direct marketing and cross-selling or up-selling; retailers, for market segmentation and store layout; and manufacturers, for quality control and maintenance scheduling. Little has been explored in the healthcare domain using data mining. Clustering can be defined as division of data into groups of similar objects without any prior knowledge. Though representation of the data objects by clusters can lead to a loss of certain finer details, simplification can be achieved at the same time. From a machine learning perspective clusters correspond to hidden patterns. The search for clusters is an example of unsupervised learning. From a practical perspective clustering plays an important role in data mining applications such as scientific data exploration, information retrieval and text mining, spatial database applications, web analysis, CRM, marketing, medical diagnostics, computational biology and many others.

Data clustering algorithms can be further sub divided into a number of categories. In this paper we concentrate on hierarchical and constraint based clustering. With hierarchical algorithms, successive clusters are found using previously established clusters, whereas constraint based clustering algorithms use some pre-defined constraints to create the clusters. Hierarchical algorithms can be agglomerate (bottom-up) or divisive (top-down). Agglomerate algorithms begin with each element as a separate cluster and merge them in successively larger clusters. The divisive algorithms begin with the whole set and proceed to divide it successively into

smaller clusters. A hierarchical algorithm in general yields a deprogram representing the nested grouping of patterns and similarity levels at which these groupings change. But there are some inherent problems associated with greedy hierarchical algorithmic approaches (AGNES, DIANA) [7] like vagueness of termination criteria of the algorithms and the inability to revisit once constructed clusters for the purpose of their improvement. To overcome these defects, an efficient technique called the multiphase clustering technique [28] can be used. This technique involves two primary phases – phase 1, where the database is scanned to built a temporary in-memory tree in order to preserve the inherent clustering structure of the data. Phase 2 selects a clustering algorithm and apply it to the leaf nodes of the tree built in the previous phase for removing sparse clusters as outliers and dense clusters into larger ones. Our approach uses this technique to build an efficient hierarchical constraint based clustering method.

Constraint based clustering is a clustering technique which deals with specific application requirements. It falls in a category in which we can divide the clustering methods into either data-driven or need-driven [3]. The data-driven clustering methods intend to discover the true structure of the underlying data by grouping similar objects together while the need-driven clustering methods group objects based on not only similarity but also needs imposed by a particular application. Thus, the clusters generated by need-driven clustering are usually more useful and actionable to meet certain application requirements. For different application needs, various balancing constraints can be designed to restrict the generated clusters and make them actionable. Particularly, we are interested in such constraint types for phase 2 of our hierarchical clustering. In this paper we will examine Gaussian distribution measures as the density based probabilistic measures which will decide the threshold for split of the leaf node clusters formed by the CF tree in phase 1. By imposing a threshold of minimum Gaussian distribution value on the clusters, our model will search for clusters which are balanced in terms of probabilistic density distribution measure. The main motivation of our paper is to develop an approach which

efficiently implements the multiphase clustering technique using Gaussian distribution measure as the constraint for creation of efficient clusters. In phase 1 first we aim at implementing a data structure similar to the CF-tree [28]. In the next phase we apply our Gaussian constraint for efficient clustering.

The main contribution of our paper are as follows:

1. We propose the multiphase clustering technique using a data structure like CF-tree and a Gaussian distribution measure as the constraint in phase 2.

2. We evaluate our algorithm using supervised clustering evaluation measure and compare it with some existing clustering techniques.

3. We also evaluate our algorithm over real data set for demonstrating the quality of the generated cluster and efficiency of the algorithm.

The remainder of the paper is organized as follows. Section II provides a list of the related works. Section III discusses in detail the proposed approach. Section IV is devoted to experimental results. We conclude our research and put forth our future works in Section V.

1 Literature Review

Clustering has been an active research area in the fields of data mining and computational statistics. Among all of the clustering methods, over the past few years, constrained (semi-supervised) based clustering methods have become very popular. Clustering with constraints is an important recent development in the clustering literature. The addition of constraints allows users to incorporate domain expertise into the clustering process by explicitly specifying what are desirable properties in a clustering solution. Constraint based clustering takes advantage of the fact that some labels (cluster membership) are known in advance, and therefore the induced

constraints are more explicit. This is particularly useful for applications in domains where considerable domain expertise already exists. It has been illustrated by several researchers that constraints can improve the results of a variety of clustering algorithms [30]. However, there can be a large variation in this improvement, even for a fixed number of constraints for a given data set. In this respect, [27] has shown that inconsistency and incoherence are two important properties of constraint based measures. These measures are strongly anti-correlated with clustering algorithm performance. Since they can be computed prior to clustering, these measures can aid in deciding which constraints to use in practice.

In real world applications such as image coding clustering, spatial clustering, geoinformatics [13], and document clustering [10], some background information is usually obtained dealing with the data objects' relationships or the approximate size of each group before conducting clustering. This information is very helpful in clustering the data [30]. However, traditional clustering algorithms do not provide effective mechanisms to make use of this information. Previous research has looked at using instance-level background information, such as must-link and cannot-link constraints [4, 26]. If two objects are known to be in the same group, we say that they are must-linked, and if they are known to be in different groups, we say that they are cannot-linked. Wagstaff *et al.* [26] incorporated this type of background information to k-means algorithm by ensuring that during the clustering process, at each iteration, each constraint is satisfied. In cannot-link constraints, the two points cannot be in the same cluster. Basu *et al.* [4] also considered must-link and cannot-link constrains to learn an underlying metric between points while clustering.

There are also works on different types of knowledge hints used in fuzzy clustering, including partial supervision where some data points have been labeled [14] and domain knowledge represented in the form of a collection of viewpoints (e.g., externally introduced prototypes/ representatives by users) [29]. Another type of constraints which has been dealt with in the

recent past is balancing constraints in which clusters are of approximately the same size or importance [3, 8]. Constraint driven clustering has been also investigated by [6], where the author proposes an algorithm based on [28] with respect to two following constraints – minimum significance constraints and minimum variance constraints. All the above mentioned researches involve specifying the number of constraint, for the formation of clusters. However, little work has been reported on constraints which make use of the Gaussian density distribution function to ascertain the probability of a data point to belong to a particular cluster.

In this paper, we aim at developing a multiphase iterative clustering algorithm using a data structure based on the CF-tree proposed by Zhang *et al.* [28] and a constraint based Gaussian density distribution function. Our algorithm will generate clusters with respect to our proposed multivariate Gaussian distribution constraint. Multivariate Gaussian distribution constraint is based on the normal density distribution of the data which can be used to ascertain the probability of that data point belonging to that cluster. The main goal of the constraint would be to check if it is possible to dynamical split the clusters formed in phase 1 of our algorithm based on an assumed threshold value so that it is advantageous for a more accurate clustering.

2 Proposed Approach

2.1 Preliminaries

Finding useful patterns in large databases without any prior knowledge has attracted a lot of research over time [28, 6]. One of widely studied problems is the identification of the clusters. There are a lot of methods which can be used for data clustering namely partitioning methods, grid based methods, hierarchical methods, density based methods, model based methods, and constraint based models [7]. The ones that we are primarily interested in this research are the hierarchical and constraint based models.

2.1.1 Centroid, Radius and Diameter

Given *n d*-dimensional data points, the centroid of the cluster is defined as the average of all points in the cluster and can be represented as

$$x_0 = \frac{\sum_{i=0}^{n} x_i}{n} \tag{1}$$

The radius of a cluster is defined as the average distance between the centroid and the member values of the cluster whereas the diameter is defined as the average distance between the individual members in the cluster

$$R = \sqrt{\frac{\sum_{i=1}^{n}(x_i - x_0)^2}{n}} \tag{2}$$

$$D = \sqrt{\frac{\sum_{i=1}^{n}\sum_{j=1}^{n}(x_i - x_j)^2}{n(n-1)}} \tag{3}$$

Although equations 1,2 ,3 are defined in [28, 7], but we redefine the concepts in this paper since they form an integral part of the paper.

2.1.2 CF-tree

In [28], the authors use an innovative data structure called the CF-tree (clustering feature) to summarize information about the data points. Clustering feature is a 3-dimensional data structure [7] containing the number of data points, their linear sum (LS) and their squared sum (SS) which can also be represented as $\sum_{i=1}^{n} x_i$ and $\sum_{i=1}^{n} x_i^2$ respectively.

$$CF = \langle n, \text{LS,SS} \rangle \tag{4}$$

2.1.3 Gaussian Distribution

In the theory of probability, Gaussian or normal distribution is defined as a continuous probability distribution of real valued random variables which tend to cluster around the centroid of a particular class.

Multivariate (k-dimensional) Gaussian distribution is a generalization of univariate normal distribution. It is generally written using the following notation: $X \sim N_k(\mu, \Sigma)$. X refers to the k-dimensional vector $(X_1, X_2, \cdots, X_k)$. μ refers to the k-dimensional mean vector and $\sum$ refers to the k x k co-variance matrix. Thus the normal distribution of a matrix where each row is an observation and each column is an attribute can be written as:

$$f_X(x) = \frac{1}{(2\pi)^{k/2}|\Sigma|^{1/2}} exp(-\frac{1}{2}(x-\mu)'\Sigma^{-1}(x-\mu)) \tag{5}$$

2.2 Overview

We start our algorithm with an input data set. Since our algorithm is best suited for continuous valued variables, we use attribute selection measure to weed out the non-continuous attributes. The resultant data set is used as the input for the first phase of our algorithm. In this phase we create a CF-tree originally proposed by [28]. The creation of CF-tree hierarchically clusters our data and all the original data points reside in the leaf nodes of the cluster. These are also referred to as the micro-clusters [7].

In the second phase of the algorithm, which is the novelty of our approach, we use the multivariate Gaussian density distribution function 5 to split the clusters. There are related research where the authors use the number of data points [28] and variance [**?**] in a cluster as measures to split a CF-Tree based cluster. Our rationale behind using the Gaussian density distribution function is, it will help us better in identifying those data points whose probability of belonging to a specific micro-cluster is low.

The Gaussian density function $f_X(x)$ varies from cluster-to-cluster. We normalize the $f_X(x)$ using the min-max normalization. Next we use a global threshold ρ to split the clusters. The data values whose $f_X(x)$ are below the threshold ρ forms a new micro-cluster.

2.3 Details

theorem From [28] - Given n d-dimensional data points in a cluster: $\{C_i\}$ where i = 1, 2, ... n, the clustering feature (CF) vector of the cluster is defined as a triple, 4.

theorem From [28] – Assume that $CF_1 = \langle N_1, LS_1, SS_1 \rangle$, $CF_2 = \langle N_2, LS_2, SS_2 \rangle$ are the CF vectors of two disjoint clusters. The CF vector of the cluster that is formed by merging the two disjoint clusters CF_1 and CF_2, is: $CF_1 + CF_2 = \langle N_1 + N_2, LS_1 + LS_2, SS_1 + SS_2 \rangle$.

2.3.1 Creation of CF-Tree

- The algorithm builds a dendrogram called clustering feature tree (CF-Tree) while scanning the data set [28].
- Each non-leaf node contains a number of child nodes. The number of children that a non-leaf node can contain is limited by a threshold called the branching factor.
- The diameter of a sub cluster under a leaf node cannot exceed a user specified threshold.
- The linear sum and squared sum for a node is already defined in 2.1.2.
- For a non-leaf node, which has child nodes $n_1, n_2, \ldots, n_k$

$$\vec{LS} = \sum_{i=1}^{k} \vec{LS} N_i \tag{6}$$

$$\vec{SS} = \sum_{i=1}^{k} \vec{SS} N_i \tag{7}$$

2.3.2 Algorithm

Phase 1: Building the CF Tree The algorithm first scans the data set and inserts the incoming data instances into the CF tree one by one. The insertion of a data instance into the CF tree is carried out by traversing the tree top-down from the root according to an instance-cluster

distance function. The data instance is inserted into the closest sub cluster under a leaf node. If the insertion of a data instance into a sub cluster causes the diameter of the sub cluster to exceed the threshold, a new sub cluster is created and is demonstrated in 1. In this figure the red and the yellow circles denote the clusters.

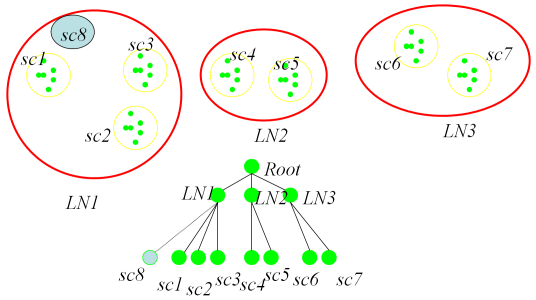

Figure 1: Insertion of a new sub cluster in the CF tree

The creation of a new leaf level sub cluster may cause its parent to cross the branching factor threshold. This causes a split in the parent node. The split of the parent node is conducted by first identifying the pair of sub clusters under the node that are separated by the largest inter-cluster distance. Then, all other sub clusters are dispatched to the two new leaf nodes based on their proximity to these two sub clusters. This is demonstrated in 2. Split of a leaf node may results in a non-leaf node containing more children than the pre-defined branching factor threshold. If so, the non-leaf nodes are split recursively based on a measure of inter-cluster distance. If the root node is split, then the height of the CF tree is increased by one, as in 3.

When the split of nodes terminates at a node, merge of the closest pair of child nodes is conducted as long as these two nodes were not formed by the latest split. A merge operation may lead to an immediate split, if the node formed by the merge contains too many child nodes. Having constructed the CF tree, the algorithm then prepares the CF tree for the phase 2 of the algorithm to cluster the micro-clusters.

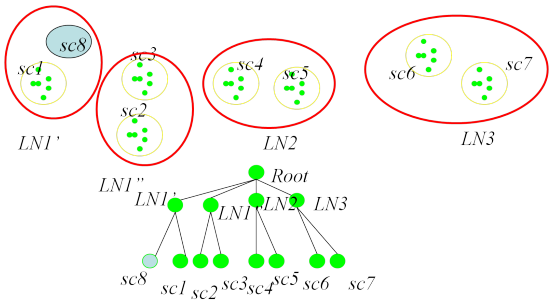

Figure 2: Split of a leaf node if it exceeds branching factor threshold

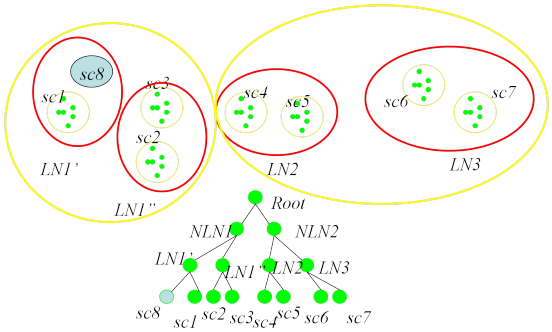

Figure 3: Split of the non-leaf nodes recursively based on inter-cluster distance measure

Phase 2: Refining using Gaussian constraint measure The CF tree obtained in phase1 now consists of clusters with sub clusters under each leaf node. Now the Gaussian constraint can be applied for the post processing of the leaf nodes of the CF tree.

We consider as a constraint the Gaussian density measure which is the threshold value used to decide the possibility of a data point in each leaf node of the CF tree to belong to that node, or in other words in that micro-cluster.

At first the Gaussian distribution function is calculated for each data point in each cluster. For this, we need two parametric inputs:

1. *Mean of all data point along each attribute*. It is represented as: $\mu_i = \sum_{i=1}^{n} x_i/n$, where n is the number of data points for each attribute in each individual clusters and x_i represents the value of the i^{th} attribute of the x^{th} data point.

2. *Covariance Matrix* is the covariance between every attribute with all other attributes in an individual clusters. Let $x = \{x_1, x_2, \ldots, x_p\}$ be a random vector with a mean of $\mu = \{\mu_1, \mu_2, \ldots, \mu_p\}$.

$$\Sigma = \begin{pmatrix} \sigma_{11} & \sigma_{12} & \ldots & \sigma_{1p} \\ \sigma_{21} & \sigma_{22} & \ldots & \sigma_{2p} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{p1} & \sigma_{p2} & \ldots & \sigma_{pp} \end{pmatrix} \tag{8}$$

The co-variance matrix is denoted as Σ. σ_{ij} is the co-variance of the i^{th}attribute with j^{th}attribute. Thus σ_{ii} is the co-variance of the i^{th} attribute with itself which are all the diagonal elements in 8. An univariate Gaussian distribution is shown in 4.

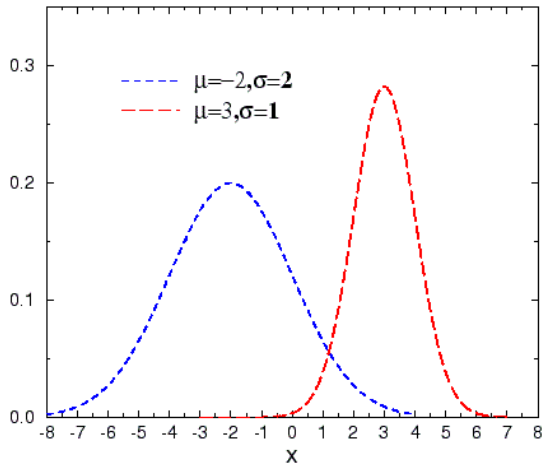

Figure 4: Univariate Gaussian distribution

The next step is to compute the Gaussian distribution function for each data point in each cluster using mean and covariance matrix calculated in the previous step.

A threshold is decided which iteratively checks all data points in each cluster to see if they fall within the threshold or not. If some data point is found to be far away from the threshold, the particular cluster to which it belongs to currently, is split. In this way all the data points are compared with the threshold. Two new clusters are formed with group of data points which

have a different range of Gaussian density measure. Hence, there is possibility of forming new sets of clusters from the leaf node of the CF tree thereby increasing its depth. The clusters formed in this way are hypothesized to be more accurate than that formed in phase 1, since we are comparing the Gaussian density measure of each data points for each cluster.

3 Experimental Results

3.1 Data Set

In order to evaluate the performance of our algorithm, we conducted a number of experiments using Abalone data set published in the year 1995 [5]. Out of the eight dimensions in the original database, we chose seven continuous dimensions (*Length*, *Diameter*, *Height*, *Whole weight*, *Shucked weight*, *Visceral weight*, *Shell weight*). The eighth dimension, *Sex*, is a nominal dimension and can take only 3 distinct values. Thus we excluded it from our experiments. In our data set we have 4177 instances each representing a single abalone inside a whole colony. The primary goal of the data set is to predict the age of an abalone from its physical measurements. Although this data set is primarily used in classification, we use it as supervised evaluation technique for cluster validity [25].

3.2 Micro-Cluster Analysis

We present the results of our micro-cluster analysis in 5. For the purposes of this test we created the initial CF-tree with varied distance measure. We started our test with a distance measure of 0.1 and went all the way to 1.0 with an increase of 0.1 in each step. The primary purpose of this test is to hypothesize on the trend of number of micro-clusters our algorithm will generate for a specific CF-tree input. The X-axis in 5 represents the distance measure and the Y-axis represents the number of micro-clusters generated. In this test we compare the number of micro-clusters generated by BIRCH [28] with our approach.

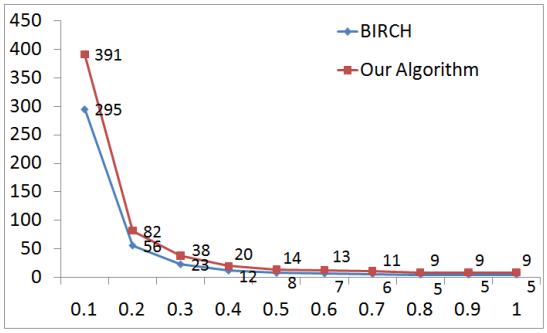

Figure 5: Micro-Cluster Analysis

From 5 we can infer that, when the number of micro-clusters generated by BIRCH is low, our algorithm generates almost twice the number of micro-clusters. But when the number of micro-clusters generated by BIRCH increases, the ratio between the number of micro-clusters generated by our algorithm and BIRCH tends to 1. The outcome of this analysis is expected, since as the number of micro-clusters generated by BIRCH increases, the amount of data points inside the micro-clusters tend to decrease significantly. When BIRCH generates a relatively large number of micro-clusters, most of the micro-clusters contain a very few data points in them which does not cross the threshold we have set in our algorithm for splitting a micro-cluster.

3.3 Scalability Test

For the scalability test we use the same data set. Since the number of tuples in the data set is not elephantine, we scale our data by appending an instance of the data set at the end of itself. In 6- X-axis denotes the number of records in our data set, whereas Y-axis represents its corresponding running times. In this way we can generate data sets which are multiples of the original data set. We present our results in figure 6.

From the results we can infer that our algorithm has a constant time component and a linearly increasing component. To lend credibility to our inference we have also provided the

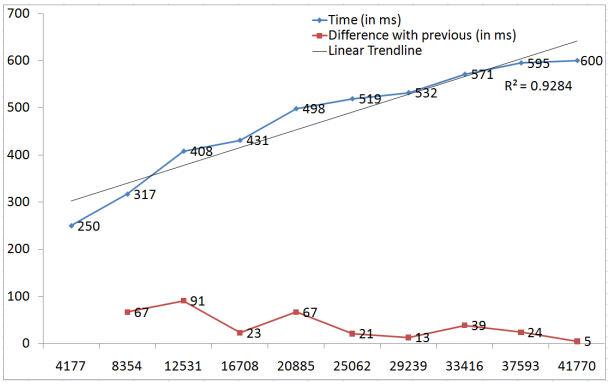

Figure 6: Scalability Test

difference in time between two consecutive results. This curve follows a horizontal straight line meaning the difference is constant.

3.4 Supervised Measures of Cluster Validity

To test the accuracy of our model we have compared our algorithm with two related approaches: BIRCH [28] and CDC [6]. We have compared our results based on the measures presented in [25] namely entropy, purity, precision and recall. For this analysis we set the distance measure in BIRCH to 0.27 so that it creates 30 micro-clusters. In response to BIRCH, our algorithm created 47 micro-clusters. We tuned the experiment in a such a way that BIRCH would create almost the same number of clusters as the number of class labels in the original data set. This provided BIRCH with a slight yet distinct advantages over our approach. The CDC algorithm was also tuned to create approximately the same number of clusters as the original class labels. To achieve this we set the significance measure in CDC to 80 and a variance of 0.34. A significance measure of 80 ensures that CDC will have between 80 to 179 nodes in all its micro-clusters. The summarized results of entropy and purity for the 3 approaches are presented in 1.

From the table we can infer that for this specific data set, our algorithm provides comparatively better results than BIRCH although our algorithm started the analysis with a slight

Table 1: Supervised Measures

	Entropy	Purity
BIRCH	2.415	0.389
Our Algorithm	2.959	0.499
CDC	3.521	0.601

disadvantage. As expected CDC out performs both the algorithms in this test. The detailed result table can be found in [1].

4 Conclusion and Future Work

In this paper we present an efficient multiphase iterative hierarchical constraint based clustering technique. Our algorithm is based on the clustering feature tree and uses Gaussian probabilistic density based distribution measure as a constraint for refining the micro-clusters. It efficiently formed clusters and ended up with a better result than the traditional BIRCH algorithm for this particular data set. Also our algorithm gives, comparable results when pitted against the CDC algorithm [6]. The scalability result shows that our algorithm performs efficiently even when tested with large data sets. We tested our algorithm and BIRCH with the same data set, and our algorithm produced more micro-clusters than BIRCH. Yet the entropy and precision ofour algorithm was higher than that of BIRCH. This indicates the accuracy of our algorithm.

We have ample scope for future work in this area. Currently our algorithm undergoes only one iteration in phase 2 for the refinement of the micro clusters. But we plan to implement iterative refinement procedure using the Gaussian density distribution constraint, for a better result. This iteration in phase 2 would continue until no more clusters can be formed. Also, we need to implement a merging algorithm which would merge the huge number of micro-clusters we hypothesize our future algorithm will form. Another important task that we consider as our future work is to test our algorithm with data sets having non continuous class level.

Acknowledgment

We would like to thank Prof. Rafal Angryk for his feedback, support and critical reviews during developing of this project.

References and Notes

[1] *http://www.cs.montana.edu/ atanu.roy/csci550/results.xlsx*.

[2] Muhammad Aurangzeb Ahmad, Brian Keegan, Atanu Roy, Dmitri Williams, Jaideep Srivastava, and Noshir S. Contractor, *Guilt by association?: network based propagation approaches for gold farmer detection*, Advances in Social Networks Analysis and Mining 2013, ASONAM '13, Niagara, ON, Canada - August 25 - 29, 2013, 2013, pp. 121–126.

[3] Arindam Banerjee, *Scalable clustering algorithms with balancing constraints*, Data Mining Knowledge Discovery **13**, 2006.

[4] Sugato Basu, A. Banjeree, ER. Mooney, Arindam Banerjee, and Raymond J. Mooney, *Active semi-supervision for pairwise constrained clustering*, In Proceedings of the 2004 SIAM International Conference on Data Mining (SDM-04, 2004, pp. 333–344.

[5] A. Frank and A. Asuncion, *UCI machine learning repository*, 2010.

[6] Rong Ge, Martin Ester, Wen Jin, and Ian Davidson, *Constraint-driven clustering*, KDD, 2007, pp. 320–329.

[7] Jiawei Han, Micheline Kamber, and Jian Pei, *Data Mining: Concepts and Techniques, Second Edition (The Morgan Kaufmann Series in Data Management Systems)*, 2 ed., Morgan Kaufmann, January 2006.

[8] Katherine A. Heller and Zoubin Ghahramani, *Bayesian hierarchical clustering*, ICML, 2005, pp. 297–304.

[9] Hian Chye Koh, Gerald Tan, et al., *Data mining applications in healthcare*, Journal of Healthcare Information ManagementVol **19** (2011), no. 2, 65.

[10] Manuel Laguia and Juan Luis Castro, *Local distance-based classification*, Know.-Based Syst. **21** (2008), 692–703.

[11] Alice Leung, William Dron, John P. Hancock, Matthew Aguirre, Jon Purnell, Jiawei Han, Chi Wang, Jaideep Srivastava, Amogh Mahapatra, Atanu Roy, and Lisa Scott, *Social patterns: Community detection using behavior-generated network datasets*, Proceedings of the 2nd IEEE Network Science Workshop, NSW 2013, April 29 - May 1, 2013, Thayer Hotel, West Point, NY, USA, 2013, pp. 82–89.

[12] Anne Milley, *Healthcare and data mining using data for clinical, customer service and financial results*, Health Management Technology **21** (2000), no. 8, 44–45.

[13] G. P. Patil, K. Sham Bhat, Raj Acharya, and S. W. Joshi, *Geoinformatic surveillance of hotspot detection, prioritization and early warning for digital governance*, DG.O, 2007, pp. 274–275.

[14] Witold Pedrycz and James Waletzky, *Fuzzy clustering with partial supervision*, IEEE Transactions on Systems, Man, and Cybernetics, Part B **27** (1997), no. 5, 787–795.

[15] Atanu Roy, Muhammad Aurangzeb Ahmad, Chandrima Sarkar, Brian Keegan, and Jaideep Srivastava, *The ones that got away: False negative estimation based approaches for gold farmer detection*, 2012 International Conference on Privacy, Security, Risk and Trust, PASSAT 2012, and 2012 International Confernece on Social Computing, SocialCom 2012, Amsterdam, Netherlands, September 3-5, 2012, 2012, pp. 328–337.

[16] Atanu Roy, Zoheb Hassan Borbora, and Jaideep Srivastava, *Socialization and trust formation: a mutual reinforcement? an exploratory analysis in an online virtual setting*, Advances in Social Networks Analysis and Mining 2013, ASONAM '13, Niagara, ON, Canada - August 25 - 29, 2013, 2013, pp. 653–660.

[17] Atanu Roy, Chandrima Sarkar, and Rafal A. Angryk, *Using taxonomies to perform aggregated querying over imprecise data*, ICDMW 2010, The 10th IEEE International Conference on Data Mining Workshops, Sydney, Australia, 13 December 2010, 2010, pp. 989–996.

[18] Chandrima Sarkar, Sumit Bhatia, Arvind Agarwal, and Juan Li, *Feature analysis for computational personality recognition using youtube personality data set*, Proceedings of the 2014 ACM Multi Media on Workshop on Computational Personality Recognition, ACM, 2014, pp. 11–14.

[19] Chandrima Sarkar, Sarah Cooley, and Jaideep Srivastava, *Robust feature selection technique using rank aggregation*, Applied Artificial Intelligence **28** (2014), no. 3, 243–257.

[20] Chandrima Sarkar, Sarah A. Cooley, and Jaideep Srivastava, *Improved feature selection for hematopoietic cell transplantation outcome prediction using rank aggregation*, Computer Science and Information Systems (FedCSIS), 2012 Federated Conference on, IEEE, 2012, pp. 221–226.

[21] Chandrima Sarkar, Prasanna Desikan, and Jaideep Srivastava, *Correlation based feature selection using rank aggregation for an improved prediction of potentially preventable events*, (2013).

[22] Chandrima Sarkar and Jaideeep Srivastava, *Predictive overlapping co-clustering*, arXiv preprint arXiv:1403.1942 (2014).

[23] Chandrima Sarkar and Jaideep Srivastava, *Impact of density of lab data in ehr for prediction of potentially preventable events*, Healthcare Informatics (ICHI), 2013 IEEE International Conference on, IEEE, 2013, pp. 529–534.

[24] Ayush Singhal, Atanu Roy, and Jaideep Srivastava, *Understanding co-evolution in large multi-relational social networks*, CoRR **abs/1407.2883** (2014).

[25] Pang-Ning Tan, Michael Steinbach, and Vipin Kumar, *Introduction to Data Mining*, us ed ed., Addison Wesley, May 2005.

[26] Kiri Wagstaff, Claire Cardie, Seth Rogers, and Stefan Schroedl, *Constrained k-means clustering with background knowledge*, In ICML, Morgan Kaufmann, 2001, pp. 577–584.

[27] Kiri L. Wagstaff, *When is constrained clustering beneficial, and why*, in AAAI, 2006.

[28] Tian Zhang, Raghu Ramakrishnan, and Miron Livny, *Birch: An efficient data clustering method for very large databases*, Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data, Montreal, Quebec, Canada, June 4-6, 1996 (H. V. Jagadish and Inderpal Singh Mumick, eds.), ACM Press, 1996, pp. 103–114.

[29] Shi Zhong and Joydeep Ghosh, *Scalable, balanced model-based clustering*, SIAM Data Mining, 2003, pp. 71–82.

[30] Shunzhi Zhu, Dingding Wang, and Tao Li, *Data clustering with size constraints*, Know.-Based Syst. **23** (2010), 883–889.